Tuning Your NT Server for the SAS® System

Jamie Walker, SAS Institute Inc.
(last revised December 1997)

Introduction

The task of optimizing or tuning Microsoft® Windows NT® for the SAS System is a logical process of determining which hardware resource is experiencing the greatest demand and then adjusting NT Server or application parameters to help relieve that demand. In most cases, the default settings provide optimal performance and do not need to be altered. However, there are cases when NT Server becomes bottlenecked and the system slows to a crawl. This paper is designed to help "fine tune" NT Server for the SAS System by outlining the four major bottlenecks and provide tips and suggestions for relieving the bottlenecks.

A bottleneck is a condition where the limitations in one of the server's major resources (memory, processor, disk I/O, network I/O) prevents the whole system from operating faster. NT Server provides tools to help monitor system performance and identify where the bottlenecks are occurring. This paper will focus on using NT Performance Monitor to help identify bottlenecks.

Performance Monitor is a graphical tool for measuring Windows NT computers. It is located in the Administrative Tools group for Windows NT Server. Performance Monitor measures the behavior of objects, such as memory, processors, physical disk, and processes, that represent components of the system. Each of these objects has associated counters that provide detailed information such as queue lengths and throughput measurements. You can use Performance Monitor to collect data over periods of time to help develop a baseline for detecting bottlenecks and to determine whether your tuning efforts are succeeding. Performance Monitor can also be used to provide charting, alerting, or reporting capabilities on current activity or ongoing logging.

The information provided in this paper is technical in nature and assumes you are familiar with the Windows NT Server environment. The rules and suggestions are provided as guidelines in detecting and resolving bottlenecks. They are not meant to be end-all solutions. You may want to refer to the *Windows NT 4.0 Workstation Resource Guide* which dedicates seven chapters to the art of performance monitoring.

Memory

Insufficient memory is by far the most common bottleneck you find when you use NT Server. A lack of RAM can cause serious performance problems. An NT Server is out of memory when the memory requirements of all active processes exceed the amount of physical memory. To solve this problem, portions of active processes are moved to disk (stored in a file named Pagefile.sys) freeing up physical memory. This process is called paging. Keeping paging at a minimum is a key to optimizing NT Server performance.

SAS System Requirements

The SAS System requirements for Windows NT recommend a minimum of 16 MB of RAM after the operating system and all other applications (running on the server) memory requirements are met. However, to achieve better performance, 32 MB or 64 MB of RAM is more ideal. When running multiple concurrent sessions of the SAS System on an NT Server, these

recommendations scale linearly. For example, if 10 concurrent SAS sessions will be running on the NT Server, a minimum of 160 MB (16 MB x 10 sessions) of RAM would be needed for the SAS System with an ideal amount anywhere from 320 MB (32 MB x 10 sessions) to 640 MB (64 MB x 10 sessions).

Performance Monitor Counters

Here are some NT performance monitor metrics to use when monitoring NT Server performance to see if you have a memory bottleneck. Use 60 second sample intervals for several hours to diagnose a memory bottleneck.

Object	Counter	Definition
Memory	Pages/sec	The number of pages read/written to disk when not found in physical RAM.
	Page Faults/sec	The total number of page faults, including hard page faults, requiring disk I/O, and soft page faults.
	Cache Faults/sec	The number of pages sought in the cache that are not present and must be found in memory or on disk.
	Available bytes	The amount of free physical memory (RAM).
Paging File	% Usage	The amount of time paging is ocurring.
	% Peak Usage	The amount of time paging is at its peak.

Rules

- If “Pages/sec” consistently averages more than 10, memory may be a bottleneck. If the average starts approaching 20, disk thrashing may be occurring.
- If “Available bytes” consistently stays below 4 MB, paging is occurring slowing down performance.
- If “Pages/sec” is greater than 10 and “Page Faults/sec” is greater than “Cache Faults/sec”, then excessive paging is taking place.

Suggestions

Properly configure the Windows NT Server service resource allocation. Using the Network applet, in the Control Panel, you can tune the NT Server service to best fit your environment. The Server Optimization Level dialog presents the user with four choices: Minimize Memory Used, Balance, Maximize Throughput for File Sharing, and Maximize Throughput for Network Applications. The following recommendations should be followed when tuning the NT Server service.

Server Optimization Level	When to Use…
“Minimize Memory Used”	1-10 concurrent SAS sessions
“Balance”	11-64 concurrent SAS sessions
“Maximize Throughput for Network Applications”	65 or more concurrent SAS sessions

The “Maximize Throughput for File Sharing” level should ONLY be used if you are using the NT Server in a file server capacity. Setting this option allows the file cache priority over applications when accessing memory. This is the default setting for NT Server. The setting you choose will affect file system and disk I/O performance.

Eliminate wasted resources. Turn off background processes and services that are unneeded by your system. Here are some examples of non-crucial services, a brief description of them, and the approximate memory they use.

Service	Description	Memory Usage
Clipbook Server	Shares clipboard files across a network.	1.3 MB
Messenger	Sends and receives pop-up messages.	49 KB
Network DDE & Network DDE DSDM	Allows dynamic data exchange across network applications.	1.6 MB
Schedule	Runs programs in the background at specified times.	1.2 MB
Telephony Server	Handles telephone dialing and remote access connections.	200 KB

NT Server creates one Pagefile.sys file, on the root drive, during installation with a default size of physical memory plus 12 MB. If possible, spread multiple paging files across multiple disk drives to take advantage of multiple disk read/write heads. Ideally, the page file(s) should be located on separate disk(s) from the application files and the Windows NT system files. As a guide to determining the size of the file on each disk, use the Performance Monitor counters "% Usage" and "% Peak Usage" of the Paging file object. Initially, you can create both with initial values equal to "% Usage" with maximum values equal to "% Peak Usage", however these values may need to be tweaked for your environment. This guarantees that no other application or process will compete with NT Server when paging needs to occur.

If you are running multiple concurrent sessions of the SAS System on your NT Server, consider using the MEMSIZE system option in the config.sas configuration file. The MEMSIZE option is used to limit the total amount of memory one session of the SAS System uses at any time. Good starting values are 32M and 64M, however, you will need experiment to find the optimal value for your system. The default value is 0 and allows the SAS System to use all available memory on the server. For more information on the MEMSIZE option, refer to the *SAS Companion for the Microsoft Windows Environment, Second Edition*.

Schedule memory-intensive processes during hours of low server activity using the AT scheduler or third party scheduling tools.

Distribute memory-intensive processes across multiple servers.

Add more memory.

Processor

A processor can execute a given number of instructions per second, largely governed by the processor's clock speed, and must be shared amongst all running processes. Processes are further subdivided into threads. Applications can be classified as single-threaded or multi-threaded. A multi-threaded application is designed to take advantage of multiple processor machines because the application threads can be distributed amongst the multiple processors. The first step to seeing whether NT Server has a processor bottleneck is to ensure the system does not have a memory bottleneck. Excessive paging or disk thrashing will cause the processor to use cycles in order to handle the paging activities.

The SAS System Architecture

The SAS System is a single-threaded application. As a result, a single invocation of the SAS System will not be able to take advantage of a multiple processor system. However, multiple, concurrent invocations of the SAS System will be distributed amongst the processors in a multiple processor system.

Performance Monitor Counters

Here are some NT performance monitor metrics to use when monitoring NT Server performance to see if you have a processor bottleneck. Use 60 second sample intervals for several hours to diagnose a processor bottleneck.

Object	Counter	Definition
Processor	% Processor Time	Percentage of time a processor is busy executing a thread other than the idle thread.
	Interrupts/sec	Number of device interrupts the processor is experiencing.
Server Work Queues	Queue Length	Length of the server work queue for the processor.

Rules

- If "% Processor Time" consistently registers at or near 100%, the processor may be the bottleneck. Next, you need to determine which process is consuming the CPU. To do this, monitor "% Processor Time" for all of the process instances.
- If "Interrupts/sec" is greater than 1000, a hardware device is generating an excessive number of interrupts. You should examine the efficiency of the hardware I/O devices such as disk controllers and network cards.

Suggestions

Adjust the priority of the job consuming the most CPU time. This can be done through the Task Manager's Process property sheet. Right-click on the process consuming the most CPU time, choose Set Priority, and then select Realtime, High, Normal, or Low. Increasing the process's priority will ensure that the process will receive more CPU time than the other running processes. The priority change takes place immediately, but is not permanent. To permanently change a process's priority, always use the start command with the specified priority from the command line to invoke the process. The priority of processes can also be controlled system-wide through the Performance property sheet in the System applet of Control Panel. The Performance property sheet allows you to change the responsiveness of foreground applications relative to background applications.

Add more processors if possible. On a multiple processor system, Windows NT will distribute the thread execution load across multiple processors. Adding processors will not improve performance if you are ONLY running one SAS session and no other applications on the server because the SAS System is a single-threaded application.

Schedule CPU-intensive processes during hours of low server activity using the AT scheduler or third party scheduling tools.

Distribute CPU-intensive processes across multiple servers.

Upgrade to a faster processor if possible.

Disk I/O

The most important aspect of I/O performance is disk performance. Disk performance not only affects I/O performance, but also directly affects virtual memory performance. For this reason, a disk bottleneck can easily be confused with a memory bottleneck.

RAID

RAID (Redundant Array of Inexpensive Disks) technology is an integral factor in disk performance optimization. There are two types of RAID, hardware and software RAID. Hardware RAID is performed by the disk controller. Software RAID is performed by the software. NT 4.0 supports hardware or software based RAID solutions or a combination of both to provide optimal performance.

RAID can be further classified into RAID levels. RAID 0, 1, and 5 are the most common hardware and software levels and are the only levels supported by NT's software RAID functions. In general, hardware based RAID solutions are faster, more reliable, offer a greater range of configuration options, and are more expensive than software based RAID solutions. The advantages of NT's built-in RAID support are the convenience of ready-made software and cost. However, using the operating system to perform RAID functions instead of offloading them to a separate hardware controller can slow server performance.

Each RAID level has different capabilities and goals. Consider the following table when choosing a RAID solution.

RAID Level	Definition	Comments
0	Disk striping	Ideal for environments where performance (read and write) is more important than fault tolerance or where maximum disk capacity is needed. RAID 0 offers no fault tolerance. If one drive in the striped set dies, the entire volume in unrecoverable.
1	Disk mirroring	Ideal for building fault-tolerant systems or data volumes without sacrificing performance. Often used for small mission critical data volumes. Disk capacity suffers significantly with RAID 1 because both drives are mirrored copies providing 50% disk capacity.
5	Disk striping with parity	Offers a reasonable compromise between performance, fault tolerance, and disk capacity. Write performance suffers significantly with RAID 5.
10	Mirrored stripe sets	Also known as RAID 0+1 or RAID 6. Must combine a hardware RAID controller with NT's RAID software capabilities. Provides the best overall performance, but the mirroring costs 50% disk capacity for all drives in the RAID 0 stripe set.

Performance Monitor Counters

Here are some NT performance monitor metrics to use when monitoring NT Server performance to see if you have a disk I/O bottleneck. To activate the disk counters, use the diskperf -ye command at the command prompt. The y option sets the system to start disk performance counters when the system is restarted. The e option allows you to measure performance of physical drives in a striped set. Use 60 second sample intervals for several hours to diagnose a disk bottleneck.

Object	Counter	Definition
Physical Disk	% Disk Time	Percentage of elapsed time the selected disk drive is busy servicing read or write requests, including wait time in the disk driver queue.
	Disk Queue Length	Pending disk I/O requests.
	Avg. Disk sec/Transfer	Time of the average disk transfer.
Memory	Pages/sec	The number of pages read/written to disk to resolve memory conflicts that were not found in RAM.

Rules

- If "% Disk Time" consistently averages near 67%, then the physical disk may be the bottleneck.
- If "Disk Queue Length" is greater than 2, then the physical disk may be the bottleneck.
- Separate the amount of disk I/O being used for paging using the following formula:
  % disk time used for paging = 100 * ("Pages/sec" * "Avg. Disk sec/Transfer")
  If the disk time used for paging is more than 10% of the total disk activity, then excessive paging is occurring.

Suggestions

Set your disk adapter's BIOS settings for maximum performance. Be sure the write-back cache has been turned on and the SCSI channel has been set to negotiate the fastest setting the disks can handle.

Use multichannel RAID controllers. Multichanneling allows use of all channels with common circuits for optimal performance, thus providing the advantages of controller duplexing without the costs.

Spread RAID volumes across channels. Use the multiple channels effectively to maximize available bandwidth. For example, put each drive in a mirrored set on its own channel or put half of a RAID 0 volume on one channel and half on another.

Use a RAID 0 stripe set for the paging file system. This will provide improved virtual memory performance.

Use a RAID 0 stripe set for the SASWORK directory. The SASWORK directory is used as a temporary storage location for files you create as well as files created internally by the SAS System. Placing SASWORK on a RAID 0 stripe set will provide improved disk I/O performance for the SAS System.

Distribute the I/O workload evenly across different disk drives. Running all applications from the root NT Server disk can cause a disk bottleneck.

Choose high-speed disk drives. Use SCSI, not IDE or EIDE. SCSI's performance far exceed IDE's for server applications. Consider the following table when choosing a drive technology.

Drive Technology	Throughput (MB/sec)
IDE	2.5
ESDI	3
SCSI - 2	5
Fast SCSI - 2	10
Fast and Wide SCSI - 2	20
SCSI - 3 or Ultra SCSI	40

SCSI - 4	80

Keep environment activity homogenous on the same device. For example, random activity and sequential activity should be on separate devices. Also, write-intensive and read-intensive environments should be on separate devices.

The performance hit for software RAID 5 is high. Use a fast RAID controller to compensate for the overhead.

Format only one logical drive per physical drive. This helps isolate disk problems, improves performance, and lowers the head movement rate over the disks.

Defragment your disks. As a file system expands and changes, contiguous blocks of space become hard to find. Performance suffers when a disk drive cannot read a file with a sequential group of operations, but must seek to retrieve the different pieces of the file. NT 4.0 does not contain a defragmentation tool. However, Executive Software markets a defrag utility called DisKeeper (http://www.diskeeper.com), as does Symantec with the Norton Utilities for NT (http://www.symantec.com).

Use the NTFS file system instead of the FAT file system. Performance does not degrade when using NTFS on large volume sizes as it does with FAT. Also, NTFS more efficiently uses disk space because of allocation differences in the two file systems. If you plan on accessing files greater than 2 gigabytes in size, you must use NTFS.

Use an online configurable RAID controller. Online configurable controllers improve manageability and uptime by providing dynamic access to your disk layout without powering down the server or taking a disk volume offline.

Properly configure the Windows NT Server service resource allocation. The file system cache competes with other applications for memory. If the disk is the bottleneck, revisit the settings suggested above for a memory bottleneck. You must decide which solution solves the most problems.

Network I/O

Network I/O becomes a system bottleneck when the demand for network resources is greater than what NT Server can provide. The network resources can be generated by clients or other server systems connecting to NT Server. Key ingredients to resolving network bottlenecks include an understanding of the network protocol and physical network you are using as well as the workloads being generated by client systems.

Performance Monitor Counters

Here are some NT performance monitor metrics to use when monitoring NT Server performance to see if you have a network bottleneck. To gain access to the network counters, install the Simple Network Management Protocol (SNMP) protocol available from the Services property sheet of the Network applet in the Control Panel. Use 60 second sample intervals for several hours to diagnose a network bottleneck.

Object	Counter	Definition
Network Interface	Output Queue Length	Pending network output packets.
	Bytes Total/sec	The number of bytes the server has sent to and received from the particular network adapter

		including all overhead incurred by the protocol in use (TCP/IP, NetBEUI, etc.) and the physical protocol (Ethernet, FDDI, etc.). This counter is dependent on the network architecture in use.
Redirector	Current Commands	Pending requests for service to the Redirector.

Rules

- If "Output Queue Length" consistently averages above 3 or 4, then the server's network adapter is waiting on the network.
- If "Bytes Total/sec" average 40 to 60 percent of the theoretical range for the network architecture in place, you may need to segment your network.

Suggestions

Distribute application usage evenly between network segments. This involves new cabling and logical subnetting, but is a proven technique and is relatively easy to implement.

Use high-speed Ethernet (100 megabits/sec) to sustain high transfer rates when moving large amounts of data across the network.

Remove unnecessary protocols not being used by your network adapter. This lowers the amount of memory NT Server allocates for network I/O and eliminates unnecessary network traffic.

Limit the number of users that can be simultaneously logged on to a server by manipulating the following Registry key:

SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\Users

Tune the Autodisconnect parameter for server connections by manipulating the following Registry key:

SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters\Autodisconnect

This parameter serves as a timer to terminate inactive connections in no open file exist and will free up server resources to accommodate active users. The default setting is 15 minutes.

If the "Current Commands" counter is much larger than the number of installed network adapter cards, then the server and/or the network is bottlenecked. You can increase the maximum number of pending network commands the redirector application I/O request queue can handle by manipulating the following Registry key:

SYSTEM\CurrentControlSet\Services\LanmanWorkstation\Parameters\MaxCmds

The default setting is 5.

Many tools, other than Performance Monitor, are available to assist in network analysis. These tools include Network Monitor (Windows NT), Unicenter TNG (Computer Associates), Open View (HP), and Works (Cisco).

Summary

Before you try to tune NT Server, establish a baseline and identify where the bottlenecks are. Fully understand your server hardware, the applications running on your system, and your network architecture. These components are interrelated and play key roles in identifying and relieving bottlenecks. Performance tuning will enable you to get the most out of the hardware you already own. However, if a hardware upgrade is your only solution, the investment you have made in performance tuning will help guide you to the right hardware upgrade solution.

Trademarks

References

Aubley, Curt (June 1997), "The Beginner's Guide to Optimizing Windows NT Server, Part 1," *Windows NT Magazine*, 123-127.

Aubley, Curt (August 1997), "The Beginner's Guide to Optimizing Windows NT Server, Part 2," *Windows NT Magazine*, 127-132.

Microsoft Corporation (1995), *Windows NT 4.0 Workstation Resource Guide*.

SAS Institute Inc. (1996), *SAS Companion for the Microsoft Windows Environment, Version 6, Second Edition*, Cary, NC: SAS Institute Inc.

Sloss, Joel (August 1997), "RAID: Enhanced Disk Storage for Windows NT," *Windows NT Magazine*, 48-56.

Suhy, Scott, "Optimization and Tuning of Windows NT," Microsoft TechNet Reference Desk, http://www.microsoft.com/syspro/technet/boes/bo/winntas/technote/nt301.htm.

Questions?

Please contact:
Joanna D'Aquanni
SAS Institute Inc.
SAS Campus Drive
Cary, NC 27513
(919) 677-8000
sasjmd@unx.sas.com

or….

Gary Mehler
SAS Institute Inc.
SAS Campus Drive
Cary, NC 27513
(919) 677-8000
gjm@unx.sas.com